arXiv:1405.0324v1 [math.AT] 1 May 2014

A sheaf-theoretic perspective on sampling

Michael Robinson

1 Introduction

Sampling theory has traditionally drawn tools from functional and complex analysis. Past successes, such as the Shannon-Nyquist theorem and recent advances in frame theory, have relied heavily on the application of geometry and analysis. The reliance on geometry and analysis means that the results are geometrically rigid.

There is evidence that topology has an important – and largely unexplored – impact on sampling problems. For instance, the space of bandlimited functions over the real line is infinite-dimensional, while the space of bandlimited functions over a compact subset is finite dimensional. There is a subtle interplay between the topology of the domain of the functions being sampled, and the class of functions themselves. For instance, one often wishes to sample from classes of non-bandlimited functions. The correct algebraic tool for modeling all of these situations is the *sheaf*; a tool which is sensitive to the topology of the domain and allows local structure to derive global inferences. For instance, bandlimited functions over manifolds can be encoded in a sheaf.

Most sampling problems that have been studied in the literature assume that samples are scalar-valued and are collected uniformly in some fashion. The theory of *sheaf morphisms* formalizes and generalizes the sampling process, allowing each sample to be vector-valued and of different dimensions.

This chapter makes several contributions to sampling theory. It proves a general sampling theorem for sheaves using the language of exact sequences. The Shannon-Nyquist theorem is a special case of this more general sampling theorem, and we show how a sheaf-theoretic approach emphasizes the impact of topology by solving several different sampling problems involving non-bandlimited functions. The solution to these problems shows that the topology of the domain has a varying level of

Michael Robinson
American University, 4400 Massachusetts Ave NW, Washington, DC 20016, e-mail: michaelr@american.edu

importance depending on the class of functions and the specific sampling question being studied.

1.1 Historical context

Sampling theory has a long and storied history, about which a number of recent survey articles [3, 13, 36, 34] have been written. Since sampling plays an important role in applications, substantial effort has been expended on practical algorithms. Our approach is topologically-motivated, like the somewhat different approach of [21, 6], so it is less constrained by specific timing constraints. Relaxed timing constraints are an important feature of bandpass [38] and multirate [37] algorithms. We focus on signals with local control, of which splines [35] are an excellent example.

Sheaf theory has not been used in applications until fairly recently. The catalyst for new applications was the technical tool of *cellular sheaves*, developed in [33]. Since that time, an applied sheaf theory literature has emerged, for instance [14, 20, 8, 27, 28].

Our sheaf-theoretic approach allows sufficient generality to treat sampling on non-Euclidean spaces. Others have studied sampling on non-Euclidean spaces, for instance general Hilbert spaces [22], Riemann surfaces [31], symmetric spaces [10], the hyperbolic plane [12], combinatorial graphs [26], and quantum graphs [23, 24]. We show that sheaves provide unified sufficiency conditions for perfect reconstruction on abstract simplicial complexes, which encompass all of the above cases.

A large class of local signals are those with *finite rate of innovation* [16, 39]. Our ambiguity sheaf is a generalization of the Strang-Fix conditions as identified in [9]. With our approach, one can additionally consider reconstruction using richer samples than simply convolutions with a function.

2 A unifying example

A celebrated consequence of the Cauchy integral formula is that the Taylor series of a holomorphic function evaluated at a point is sufficient to determine its value anywhere in its domain, if this is connected. Analytic continuation is therefore a very strong kind of reconstruction from a single sample. Analytic continuation relies both on (1) a restricted space of functions (merely smooth functions do not suffice) and (2) a rather large amount of information at the sample point (not just the value of the function, but also all of its derivatives). These two constraints are essential to understand the nature of reconstruction from samples, so the admittedly special case of analytic continuation is informative.

Consider the space of holomorphic functions $C^\omega(U,\mathbb{C})$ on a connected open set $U \subseteq \mathbb{C}$. Without loss of generality, suppose that U contains the origin. Then the function $a : C^\omega(U,\mathbb{C}) \to l^1$ given by

$$a(f) = \left(f(0), f'(0), \ldots, \frac{f^{(n)}(0)}{n!}, \ldots \right)$$

for $f \in C^\omega(U, \mathbb{C})$ is a linear transformation. Because a computes the Taylor series of f, whenever $a(f) = a(g)$ it must follow that $f = g$ on U. This means that as a linear transformation, the sampling function a has a trivial kernel.

Conversely, the trivial kernel of a witnesses the fact that the original $f \in C^\omega(U, \mathbb{C})$ can be recovered from the sampled value $a(f)$. This is by no means necessarily true for all sampling functions. For instance, the sampling function $b : C^\omega(U, \mathbb{C}) \to l^1$ given by

$$b(f) = \left(0, f'(0), \ldots, \frac{f^{(n)}(0)}{n!}, \ldots \right)$$

has a one-dimensional kernel. This means that reconstruction of an analytic function from its image through b is ambiguous – it is known only up to the addition of a constant. But there is more information available than merely the *dimension* of $\ker b$, since it is a subspace of $C^\omega(U, \mathbb{C})$. Indeed, if we restrict the domain of b to be the subspace $Z \subset C^\omega(U, \mathbb{C})$ of analytic functions whose value at the origin is zero, then the intersection $Z \cap \ker b$ is trivial. Reconstruction succeeds on Z using b for sampling even though using b on its whole domain is ambiguous.

Sampling a function in $C^\omega(U, \mathbb{C})$ can be represented generally as a function $s : C^\omega(U, \mathbb{C}) \to l^1$. Observe that s could take the form of the functions a or b above, in which a function is evaluated in the immediate vicinity of a single point. However, s could also be given by

$$s(f) = (\ldots, f(-1), f(0), f(1), \ldots)$$

or many other possibilities. In this general setting, the simplest way to determine whether reconstruction is ambiguous is to examine $\ker s$. Recognizing that we may wish to restrict the class of functions under discussion, it is useful to understand how the subspace $\ker s$ is included within $C^\omega(U, \mathbb{C})$. This situation can also be described as the following *exact sequence* of linear functions

$$0 \to A \xrightarrow{\ i\ } C^\omega(U, \mathbb{C}) \xrightarrow{\ s\ } l^1 \to 0,$$

which means that $\ker s = \operatorname{image} i$. Observe that the zero at the beginning of the sequence indicates that the map i is injective, so that $A = \ker s$. Likewise, the zero at the end of the sequence indicates that s is surjective.

If $f \in C^\omega(U, \mathbb{C})$ is fixed (but unknown) and $s(f)$ is known, then clearly f can only be known to be one of the preimages $i^{-1}(f) \subseteq A$. If there is only one preimage (as in the case of a above), then reconstruction is said to be *unambiguous*.

This example contrasts sharply with the situation of sampling data from the space of *all* smooth functions. In this case, one has a diagram like

$$0 \to C \xrightarrow{i} C^\infty(U,\mathbb{C}) \xrightarrow{c} l^1 \to 0,$$

because in this case C is quite large. The analytic functions are a subset $C^\omega(U,\mathbb{C}) \subset C^\infty(U,\mathbb{C})$. This can be expressed diagrammatically as

$$\begin{array}{ccccccccc}
 & & & & 0 & & 0 & & \\
 & & & & \downarrow & & \downarrow & & \\
 & & 0 & \xrightarrow{i} & C^\omega(U,\mathbb{C}) & \xrightarrow{a} & l^1 & \longrightarrow & 0 \\
 & & \downarrow & & \downarrow & & \downarrow & & \\
0 & \longrightarrow & C & \xrightarrow{i} & C^\infty(U,\mathbb{C}) & \xrightarrow{c} & l^1 & \longrightarrow & 0
\end{array}$$

in which every possible composition of linear maps with the same domain and codomain are equal. This shows how the two classes of functions and their samples are related, and the technique will be used in later sections as a kind of algebaic bound.

3 Local data

Vector spaces of functions such as $C^k(U,\mathbb{C})$ are rather global in nature – an element of such a space *is* a function! In contrast, evaluating a function at a particular point x corresponds to a linear transformation that is only sensitive to a function's value at or nearby x. Because function evaluation is a local process, reconstructing the global function space element from these samples appears counterintuitive.

The local sampling versus global reconstruction paradox is resolved because reconstruction theorems only exist for certain suitably constrained vector spaces. For instance, the Paley-Wiener space PW_B consists of functions f whose Fourier transform

$$\hat{f}(\omega) = \int_{-\infty}^{\infty} f(x) e^{-2\pi i \omega x} dx$$

is supported on $[-B,B]$. We say that each $f \in PW_B$ has *bandwidth* B. The Shannon-Nyquist theorem asserts that functions in $PW_{1/2}$ are uniquely determined by their values on the integers, which is best explained by the fact that every $f \in PW_{1/2}$ has a cardinal series decomposition

$$f(x) = \sum_{n=-\infty}^{\infty} f(n) \frac{\sin \pi(x-n)}{\pi(x-n)}$$

where each *sinc function* given by $\operatorname{sinc}(x-n) = \frac{\sin(x-n)}{x-n}$ has bandwidth $1/2$ or less. Moreover, the set of sinc functions is orthonormal over the usual inner product in $PW_{1/2}$, so we have that

$$f(x) = \int_{n=-\infty}^{\infty} f(n) \frac{\sin \pi(x-n)}{\pi(x-n)}\, dx. \tag{1}$$

Even though the support of $\operatorname{sinc} \pi(x-n)$ is $\mathbb{R}$, it decays away from n. This means that in (1), the effect of values of f far away from n will have little effect on $f(n)$. So in the case of PW_B, sampling via (1) is only *approximately* local. Because of this, global constraints – such as those arising from compactness – on the function space play an important role in sampling theorems.

This section formalizes the above intuition, by constructing a *sheaf theoretic* framework for discussing sampling. Sheaves are the correct mathematical formalism for discussing local information. Section 3.1 distills an axiomatic framework that precisely characterizes what "local" means. Section 3.2 defines the cohomology functor for sheaves, which assembles this local information into global information. With the definition of a *sheaf morphism* in Section 3.3, these tools allow the statement of general conditions under which a sampling suffices to reconstruct a function in a particular space in Section 4.

3.1 Sheaves represent local data

A local model of data should be flexible enough to capture both analytic and non-analytic functions. Because portions of the data in one region will not necessarily be related to those farther away, the model should allow us to infer global effects only when they are appropriate to the kind of function under study.

Spaces of continuous functions exhibit several properties related to locality. As a concrete example, consider the following properties of $C^k(U,\mathbb{C})$ when $k \geq 0$:

1. *Restriction*: Whenever $V \subseteq U$ are open sets, there is a linear map $C^k(U,\mathbb{C}) \to C^k(V,\mathbb{C})$ that is given by restricting the domain of a function defined on U to one defined on V.
2. *Uniqueness*: Whenever a function is the zero function on some open set, then all of its restrictions are zero functions also. The converse is true also: suppose $f \in C^k(V,\mathbb{C})$, and that $\{U_1,\dots\}$ is an open cover of V. If the restriction of function f to each U_k is the zero function on U_k, then f has to be the zero function on V.
3. *Gluing*: If U and V are open sets and $f \in C^k(U,\mathbb{C})$, $g \in C^k(V,\mathbb{C})$ then whenever $f(x) = g(x)$ for all $x \in U \cap V$ there is a function $h \in C^k(U \cup V,\mathbb{C})$ that restricts to f and g.

The gluing property provides a condition by which local information (the elements $f \in C^k(U,\mathbb{C})$, $g \in C^k(V,\mathbb{C})$) can be assembled into global information in $C^k(U \cup V,\mathbb{C})$, provided a consistency condition is met. We will call this specification of f and g a *section* when they restrict to the same element in $C^k(U \cap V,\mathbb{C})$. This can also be illustrated diagrammatically

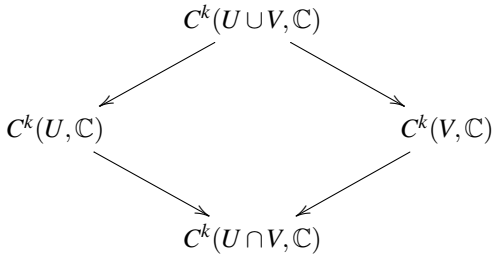

where the arrows represent the restrictions of functions from one domain to the next. Specifically, when two functions on the middle level are mapped to the same function on the bottom level, they are both images of a function on the top level.

Let's formalize these properties to obtain a more general construction, in which the data are not necessarily encoded as continuous functions. It is usually unnecessary to consider *all* open sets; what's really relevant is the intersection lattice. In this chapter, we need a concept of space that is convenient for computations. The most efficient such definition is that of a simplicial complex.

Definition 1. An *abstract simplicial complex* X on a set A is a collection of ordered subsets of A that is closed under the operation of taking subsets. We call each element of X a *face*. A face with $k+1$ elements is called a k-dimensional face (or a k-face), though we usually call a 0-face a *vertex* and a 1-face an *edge*. If all of the faces of an abstract simplicial complex X are of dimension n or less, we say that X is an n-dimensional simplicial complex. If X is a 1-dimensional simplicial complex, we usually call X a *graph*.

The *face category* has the elements of X for its objects, and setwise inclusions of one element of X into another for its morphisms. If a and b are two faces in an abstract simplicial complex X with $a \subset b$ and $|a| < |b|$, we will write $a \rightsquigarrow b$ and say that a is *attached* to b. Finally, a collection Y of faces of X is called a *closed subcomplex* if whenever $b \in Y$ and $a \rightsquigarrow b$, then $a \in Y$ also.

Sometimes simplicial complexes arise naturally from the problem, for instance the connection graph for a network, but it is helpful to have a procedure to obtain a simplicial complex from a topological space. Suppose that X is a topological space and that $\mathscr{U} = \{U_1, \dots\}$ is an open cover of X.

Definition 2. The *nerve* $N(\mathscr{U})$ is the abstract simplicial complex whose vertices are given by the elements of $\mathscr{U}$, and whose k-faces $\{U_{i_0}, \dots, U_{i_k}\}$ are given by the nonempty intersections $U_{i_0} \cap \cdots \cap U_{i_k}$.

Example 1. Figure 1 shows two covers and their associated nerves. In the left diagram, the sets A, B, and C have nonempty pairwise intersections and a nonempty triple intersection $A \cap B \cap C$, so the nerve is a 2-dimensional abstract simplicial complex. In the right diagram, $A \cap B \cap C$ is empty, so the nerve is only 1-dimensional.

The concept of local information over a simplicial complex is a straightforward generalization of the three properties (restriction, uniqueness, and gluing) for continuous functions. The resulting mathematical object is called a *sheaf*.

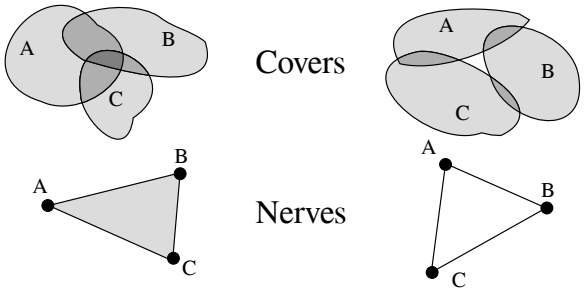

Fig. 1 The nerve of two covers: (left) with a nonempty triple intersection (right) without a triple intersection

Definition 3. A *sheaf* $\mathscr{F}$ on an abstract simplicial complex X is a covariant functor from the face category of X to the category of vector spaces. Explicitly,

- for each element a of X, $\mathscr{F}(a)$ is a vector space, called the *stalk at* a,
- for each attachment of two faces $a \rightsquigarrow b$ of X, $\mathscr{F}(a \rightsquigarrow b)$ is a linear function from $\mathscr{F}(a) \to \mathscr{F}(b)$ called a *restriction*, and
- for every composition of attachments $a \rightsquigarrow b \rightsquigarrow c$, the restrictions satisfy $\mathscr{F}(b \rightsquigarrow c) \circ \mathscr{F}(a \rightsquigarrow b) = \mathscr{F}(a \rightsquigarrow b \rightsquigarrow c)$.

We will usually refer to X as the *base space* for $\mathscr{F}$.

Remark 1. Although sheaves have been extensively studied over topological spaces (see [4] or the appendix of [18] for a modern, standard treatment), the resulting definition is ill-suited for application to sampling. Instead, we follow a substantially more combinatorial approach introduced in the 1980 thesis of Shepard [32].

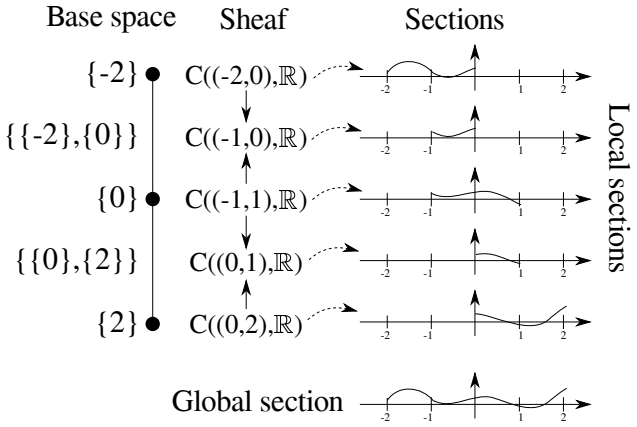

Fig. 2 A sheaf of continuous functions over an interval (compare with Figure 3)

Example 2. The space of continuous functions over a topological space can be represented as a sheaf. For instance, Figure 2 shows one way to organize the space of continuous functions over the interval $(-2,2)$ in terms of spaces of continuous functions over smaller intervals. (See Example 5 for another encoding of continuous functions as a sheaf.) In this particular sheaf model, the base space is given by an abstract simplicial complex X over three abstract vertices, which we label suggestively as $\{-2\}$, $\{0\}$, $\{2\}$. The simplicial complex X has two edges as shown in the diagram. We define the sheaf $\mathscr{C}$ over X by assigning spaces of continuous functions to each face, and define the restrictions between stalks to be the process of "actually" restricting the functions.

Example 3. Coming back to Section 2, a sheaf of analytic functions can be constructed as a subsheaf of the previous example by merely replacing the stalks with spaces of analytic functions defined over the appropriate intervals.

Notice that the definition of a sheaf captures the *restriction* locality property, but does not formalize the *uniqueness* or *gluing* properties. Some authors [4, 7, 19, 15] explicitly require these properties from the outset, calling the object defined in Definition 3 a *presheaf*, regarding it as incomplete. Although the difference between sheaves and presheaves is useful in navigating certain technical arguments, every presheaf has a unique *sheafification*. Because of this, our strategy follows the somewhat more economical treatment set forth in [32, 30], which removes this distinction. As a consequence of this choice, we explicitly define collections of *sections*, which effectively implement the sheafification.

Definition 4. Suppose $\mathscr{F}$ is a sheaf on an abstract simplicial complex X and that $\mathscr{U}$ is a collection of faces of X. An assignment s which assigns an element of $\mathscr{F}(a)$ to each face $a \in \mathscr{U}$ is called a *section* supported on $\mathscr{U}$ when for each attachment $a \rightsquigarrow b$ of faces in $\mathscr{U}$, $\mathscr{F}(a \rightsquigarrow b)s(a) = s(b)$. We will denote the space of sections of $\mathscr{F}$ over $\mathscr{U}$ by $\mathscr{F}(\mathscr{U})$, which is easily checked to be a vector space. A *global section* is a section supported on X. If r and s are sections supported on $\mathscr{U} \subset \mathscr{V}$, respectively, in which $r(a) = s(a)$ for each $a \in \mathscr{U}$ we say that *s extends r*.

Example 4. Consider $Y \subseteq X$ a subset of the vertices of an abstract simplicial complex. The sheaf $\mathscr{S}$ which assigns a vector space V to vertices in Y and the trivial vector space to every other face is called a *V-sampling sheaf supported on Y*. To every attachment of faces of different dimension, S will assign the zero function. For a finite abstract simplicial complex X, the space of global sections of a V-sampling sheaf supported on Y is isomorphic to $\bigoplus_{y \in Y} V$.

Example 5. Figure 5 shows a sheaf whose global sections are continuous functions that is essentially dual to the one in Example 2. (Although it is straightforward to generalize the construction to cell complexes of arbitrary dimension, we will work over an interval to keep the exposition simple.) Specifically, consider a simplicial complex with two vertices v_1 and v_2 and one edge e between them. The stalk over each vertex is a space of continuous functions as in Example 2, though we require

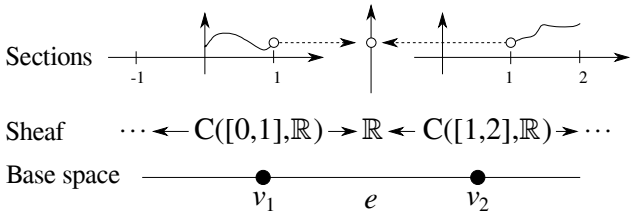

Fig. 3 Another sheaf of continuous functions over an interval (compare with Figure 2)

the functions to be continuous over a *closed* interval. However, the stalk over the edge is merely $\mathbb{R}$. The restriction in this case evaluates functions at an appropriate endpoint. If we name the sheaf $\mathscr{C}$, then for $f \in \mathscr{C}(v_1) = C([0,1],\mathbb{R})$,

$$(\mathscr{C}(v_1 \rightsquigarrow e))(f) = f(1),$$

and

$$(\mathscr{C}(v_2 \rightsquigarrow e))(g) = g(1),$$

for $g \in \mathscr{C}(v_2) = C([1,2],\mathbb{R})$. Observe that the global sections of this sheaf are precisely functions that are continuous on $[0,2]$.

This second example of a sheaf of continuous function can be easily recast to describe discrete timeseries as well. However, it is convenient to discriminate between the restrictions "to the left" versus those "to the right." This is most conveniently described by the concept of orientation. Recall that an abstract simplicial complex X consists of *ordered* sets. For a a k-face and b a $k+1$-face, define the *orientation index*

$$[b:a] = \begin{cases} +1 & \text{if the order of elements in } a \text{ and } b \text{ agrees,} \\ -1 & \text{if it disagrees, or} \\ 0 & \text{if } a \text{ is not a face of } b. \end{cases}$$

Example 6. If V is a vector space, then the m-term grouping sheaf $\mathscr{V}^{(m)}$ has the diagram

$$\xrightarrow{\sigma_+} V^{m-1} \underset{\sigma_-}{\longleftarrow} V^m \xrightarrow{\sigma_+} V^{m-1} \underset{\sigma_-}{\longleftarrow}$$

in which the restrictions are given by

$$\sigma_-(x_1,\ldots,x_m) = (x_1,\ldots,x_{m-1}) \text{ and } \sigma_+(x_1,\ldots,x_m) = (x_2,\ldots,x_m).$$

Observe that the sampling sheaf defined in Example 4 is merely $V^{(1)}$, and that the space of global sections of all grouping sheaves are isomorphic.

Sheaves can also describe spaces of piecewise continuous functions, as the next example shows.

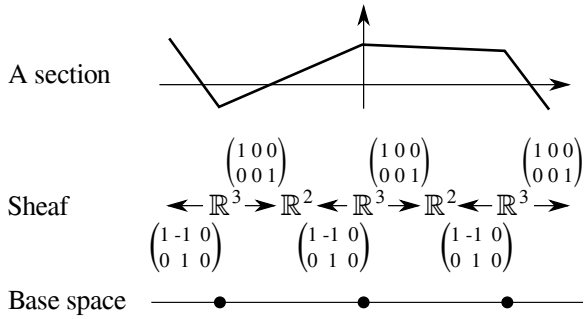

Fig. 4 An example of a sheaf $\mathscr{PL}$ over a graph

Example 7. Suppose G is a graph in which each vertex has finite degree. Let $\mathscr{PL}$ be the sheaf constructed on G that assigns $\mathscr{PL}(v) = \mathbb{R}^{1+\deg v}$ to each edge v of degree $\deg v$ and $\mathscr{PL}(e) = \mathbb{R}^2$ to each edge e. The stalks of $\mathscr{PL}$ specify the value of the function (denoted y below) at each face and the slopes of the function on the edges (denoted $m_1, \ldots, m_k$ below). To each attachment of a degree k vertex v into an edge e, let $\mathscr{PL}$ assign the linear function

$$(\mathscr{PL}(v \rightsquigarrow e))(y, m_1, \ldots, m_e, \ldots, m_k) = \begin{pmatrix} y + ([e:v]-1)\frac{1}{2}m_e \\ m_e \end{pmatrix}.$$

The global sections of this sheaf are *piecewise linear functions* on G; see Figure 4.

3.2 Sheaf cohomology

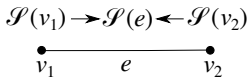

Fig. 5 A sheaf over a small abstract simplicial complex

The space of global sections of a sheaf is important in applications. Although Definition 4 is not constructive, one can compute this space algorithmically. Specifically, consider the abstract simplicial complex X shown in Figure 5, which consists of an edge e between two vertices v_1 and v_2. Suppose s is a global section of a sheaf $\mathscr{S}$ on X. This means that

$$\mathscr{S}(v_1 \rightsquigarrow e)s(v_1) = s(e) = \mathscr{S}(v_2 \rightsquigarrow e)s(v_2).$$

Since the above equation is written in a vector space, we can rearrange it to obtain the equivalent specification

$$\mathscr{S}(v_1 \rightsquigarrow e)s(v_1) - \mathscr{S}(v_2 \rightsquigarrow e)s(v_2) = 0,$$

which could be written in matrix form as

$$\begin{pmatrix}\mathscr{S}(v_1 \rightsquigarrow e) & -\mathscr{S}(v_2 \rightsquigarrow e)\end{pmatrix}\begin{pmatrix}s(v_1)\\ s(v_2)\end{pmatrix} = 0.$$

This purely algebraic manipulation shows that computing the space of global sections of a sheaf is equivalent to computing the kernel of a particular matrix as in Section 2. Clearly this procedure ought to work for arbitrary sheaves over arbitrary abstract simplicial complexes, though it could get quite complicated. *Cohomology* is a systematic way to perform this computation, and it results in additional information as we'll see in later sections.

The vector $\begin{pmatrix}s(v_1)\\ s(v_2)\end{pmatrix}$ above suggests that we should define following formal *cochain* vector spaces $C^k(X;\mathscr{F}) = \bigoplus_{a \text{ a } k\text{-face of } X} \mathscr{F}(a)$ to represent the possible choices of data over the k-faces. In the same way, the matrix

$$\begin{pmatrix}\mathscr{S}(v_1 \rightsquigarrow e) & -\mathscr{S}(v_2 \rightsquigarrow e)\end{pmatrix}$$

generalizes into the *coboundary map* $d^k : C^k(X;\mathscr{F}) \to C^{k+1}(X;\mathscr{F})$, which we now define. The coboundary map d^k takes an assignment s on the k-faces to a different assignment $d^k s$ whose value at a $(k+1)$-face b is

$$(d^k s)(b) = \sum_{a \text{ a } k\text{-face of} X} [b : a]\mathscr{F}(a \rightsquigarrow b)s(a).$$

Together, we have a sequence of linear maps

$$0 \to C^0(X;\mathscr{F}) \xrightarrow{d^0} C^1(X;\mathscr{F}) \xrightarrow{d^1} C^2(X;\mathscr{F}) \xrightarrow{d^2} \cdots$$

called the *cochain complex*.

As in the simple example described above, the kernel of d^k consists of data specified on k-faces that is consistent, when tested on the $(k+1)$-faces. However, it can be shown that $d^{k+1} \circ d^k = 0$, so that the image of d^k is a subspace of the kernel of d^{k+1}. This means that the image of d^k is essentially redundant information, since it is already known to be consistent when tested on the $(k+2)$-faces. Because of this fact, only those elements of the kernel of d^k that are *not* already known to be consistent are really worth mentioning. This leads to the definition of sheaf cohomology:

Definition 5. The k-th *sheaf cohomology* of $\mathscr{F}$ on an abstract simplicial complex X is

$$H^k(X;\mathscr{F}) = \ker d^k / \operatorname{image} d^{k-1}.$$

As an immediate consequence of this construction, we have the following useful statement.

Proposition 1. *$H^0(X;\mathscr{F}) = \ker d^0$ consists precisely of those assignments s which are global sections, so a global section is determined entirely by its values on the vertices of X.*

3.3 Transformations of local data

Sheaves can be used to represent local data and cohomology can be used to infer the resulting globally-consistent data. However interesting this theory may be, we need to connect it to the process of sampling. Indeed, as envisioned in Section 2, sampling is a *transformation* between two spaces of functions – from functions with a continuous domain to functions with a discrete domain. Such a transformation arising from sampling respects the local structure of the function spaces. This kind of transformation is called a *sheaf morphism*. There are two aspects to a sheaf morphism: (1) its effect on the base space, and (2) its effect on stalks. The effect on the base space should be to respect local neighborhoods, which means that a sheaf morphism must at least specify a continuous map. Since we have restricted our attention to abstract simplicial complexes rather than general topological spaces, the analog of a continuous map is a simplicial map.

Definition 6. A *simplicial map* from one abstract simplicial complex X to another Y is a function f from the set of simplices of X to the simplices of Y that additionally satisfies two properties:

1. If $a \rightsquigarrow b$ is an attachment of two simplices in X, then $f(a) \rightsquigarrow f(b)$ is an attachment of simplices in Y, and
2. The dimension of $f(a)$ is no more than the dimension of a, a simplex in X.

The last condition means a simplicial map takes vertices to vertices, edges either to edges or vertices, and so on.

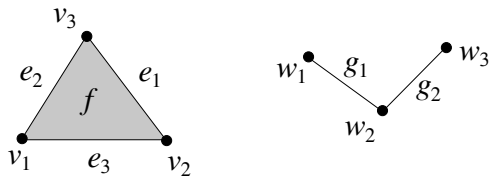

Fig. 6 The simplicial complexes X (left) and Y (right) for Example 8

Example 8. Consider the simplicial complexes X and Y shown in Figure 6. The function $F : X \to Y$ given by

$$F(v_1) = w_1,\ F(v_2) = w_2,\ F(v_3) = w_2$$

determines a simplicial map, in which $F(e_1) = w_2$, $F(e_2) = F(e_3) = F(f) = g_1$.

In contrast, any function that takes v_1 to w_1, v_2 to w_2, and v_3 to w_3 cannot be a simplicial map because the image of e_2 should be an edge from w_1 to w_3, but no such edge exists.

Definition 7. Suppose that $f : X \to Y$ is a simplicial map, and that $\mathscr{F}$ is a sheaf on Y and $\mathscr{G}$ is a sheaf on X. A *sheaf morphism* (or simply a *morphism*) $m : \mathscr{F} \to \mathscr{G}$ along f assigns a linear map $m_a : \mathscr{F}(f(a)) \to \mathscr{G}(a)$ to each face $a \in X$ so that for every attachment $a \rightsquigarrow b$ in the face category of X, $m_b \circ \mathscr{F}(f(a) \rightsquigarrow f(b)) = \mathscr{G}(a \rightsquigarrow b) \circ m_a$.

Usually, we describe a morphism by way of a commutative diagram like the one below

$$\begin{array}{ccc} \mathscr{F}(f(a)) & \xrightarrow{m_a} & \mathscr{G}(a) \\ {\scriptstyle \mathscr{F}(f(a) \rightsquigarrow f(b))}\downarrow & & \downarrow{\scriptstyle \mathscr{G}(a \rightsquigarrow b)} \\ \mathscr{F}(f(b)) & \xrightarrow{m_b} & \mathscr{G}(b) \end{array}$$

Remark 2. The reader is cautioned that a sheaf morphism and its underlying simplicial map "go opposite ways."

Cohomology is a functor from the category of sheaves and sheaf morphisms to the category of vector spaces. This indicates that cohomology preserves and reflects the underlying relationships between data stored in sheaves.

Proposition 2. *Suppose that $\mathscr{R}$ is a sheaf on X and that $\mathscr{S}$ is a sheaf on Y. If $m : \mathscr{R} \to \mathscr{S}$ is a morphism of these sheaves, then m induces linear maps $m^k : H^k(X, \mathscr{R}) \to H^k(Y, \mathscr{S})$ for each k. (Note that the simplicial map associated to m is a function $Y \to X$.)*

As a consequence, m^0 is a linear map from the space of global sections of $\mathscr{R}$ to the space of global sections of $\mathscr{S}$. Because of this, it is possible to describe the process of sampling using a sheaf morphism.

Definition 8. Suppose that $\mathscr{F}$ is a sheaf on an abstract simplicial complex X, and that $\mathscr{S}$ is a V-sampling sheaf on X supported on a closed subcomplex Y. A *sampling morphism* (or *sampling*) of $\mathscr{F}$ is a morphism $s : \mathscr{F} \to \mathscr{S}$ that is surjective on every stalk.

Example 9. The diagram below shows a morphism (vertical arrows) between two sheaves, namely the sheaf of continuous functions defined in Example 2 (top row) and the sampling sheaf defined in Example 4 (bottom row):

$$\begin{array}{ccccccc} C((-2,0),\mathbb{R}) & \longrightarrow & C((-1,0),\mathbb{R}) & \longleftarrow & C((-1,1),\mathbb{R}) & \longrightarrow & C((0,1),\mathbb{R}) \\ \downarrow{\scriptstyle e_{-1}} & & \downarrow & & \downarrow{\scriptstyle e_0} & & \downarrow \\ \mathbb{R} & \longrightarrow & 0 & \longleftarrow & \mathbb{R} & \longrightarrow & 0 \end{array}$$

In the diagram, e_x represents the operation of evaluating a continuous function at x. As in Section 2, this sampling morphism takes a continuous function $f \in C((-2,1),\mathbb{R})$ to a vector $(f(-1),f(0))$.

In algebraic topology, special emphasis is placed on *sequences* of maps of the form

$$\cdots \longrightarrow A_1 \xrightarrow{m_1} A_2 \xrightarrow{m_2} A_3 \xrightarrow{m_3} A_4 \xrightarrow{m_4} \cdots,$$

where the A_k are vector spaces and the m_k are linear maps. We will denote this sequence by $(A_\bullet, m_\bullet)$. For instance, the cochain complex described in the previous section is a sequence of vector spaces. A linear map satisfies the dimension theorem, which relates the size of its kernel, cokernel, and image. In some sequences, the dimension theorem is extremely useful – these are the exact sequences.

Definition 9. A sequence $(A_\bullet, m_\bullet)$ of vector spaces is called *exact* if $\ker m_k = \operatorname{image} m_{k-1}$.

Via the dimension theorem, exact sequences can encode information about linear maps, namely

1. $0 \to A \xrightarrow{m} B$ is exact if and only if m is injective,
2. $A \xrightarrow{m} B \to 0$ is exact if and only if m is surjective, and
3. $0 \to A \xrightarrow{m} B \to$ is exact if and only if m is an isomorphism.

Observe that the cochain complex $(C^\bullet(X;\mathscr{S}),d^\bullet)$ is exact if and only if $H^k(X;\mathscr{S}) = 0$ for all k.

Remark 3. Sequences of sheaf morphisms (instead of just vector spaces) are surprisingly powerful, and play an important role in the general theory of sheaves. However, if the direction of the morphisms is allowed to change across the sequence, like

$$\mathscr{A} \longleftarrow \mathscr{B} \longrightarrow \mathscr{C},$$

the resulting construction can represent all linear, shift-invariant filters [29, 30].

4 The general sampling theorems

Given a sampling morphism, we can construct the *ambiguity sheaf* $\mathscr{A}$ in which the stalk $\mathscr{A}(a)$ for a face $a \in X$ is given by the kernel of the map $\mathscr{F}(a) \to \mathscr{S}(a)$. If $a \rightsquigarrow b$ is an attachment of faces in X, then $\mathscr{A}(a \rightsquigarrow b)$ is given by $\mathscr{F}(a \rightsquigarrow b)$ restricted to $\mathscr{A}(a)$. This implies that the sequence of sheaves

$$0 \to \mathscr{A} \longrightarrow \mathscr{F} \xrightarrow{s} \mathscr{S} \to 0$$

induces short exact sequences of cochain spaces

$$0 \to C^k(X;\mathscr{A}) \to C^k(X;\mathscr{F}) \to C^k(X;\mathscr{S}) \to 0,$$

one for each k. Together, these sequences of cochain spaces induce a long exact sequence (via the well-known Snake lemma; see [17] for instance)

$$0 \to H^0(X;\mathscr{A}) \to H^0(X;\mathscr{F}) \to H^0(X;\mathscr{S}) \to H^1(X;\mathscr{A}) \to \cdots$$

An immediate consequence is therefore

Corollary 1. *(Sheaf-theoretic Nyquist theorem) The global sections of $\mathscr{F}$ are identical with the global sections of $\mathscr{S}$ if and only if $H^k(X;\mathscr{A}) = 0$ for $k = 0$ and 1.*

The cohomology space $H^0(X;\mathscr{A})$ characterizes the *ambiguity* in the sampling. When $H^0(X;\mathscr{A}$ is nontrivial, there are multiple global sections of $\mathscr{F}$ that result in the same set of samples. In contrast, $H^1(X;\mathscr{A})$ characterizes the *redundancy* of the sampling. When $H^1(X;\mathscr{A}$ is nontrivial, then there are sets of samples that correspond to *no* global section of $\mathscr{F}$. Optimal sampling therefore consists of identifying minimal closed subcomplexes Y so the resulting ambiguity sheaf $\mathscr{A}$ has $H^0(X;\mathscr{A}) = H^1(X;\mathscr{A}) = 0$.

Remark 4. Corollary 1 is also useful for describing boundary value problems for differential equations. The sheaf $\mathscr{F}$ can be taken to be a sheaf of solutions to a differential equation [11]. The sheaf $\mathscr{S}$ can be taken to have support only at the boundary of the region of interest, and therefore specifies the possible boundary conditions. In this case, the space of global sections of the ambiguity sheaf $\mathscr{A}$ consists of all solutions to the differential equation that *also* satisfy the boundary conditions.

Let us place bounds on the cohomologies of the ambiguity sheaf. To do so, we construct two new sheaves associated to a given sheaf $\mathscr{F}$ and a closed subcomplex $Y \subseteq X$. These new sheaves allow us to study reconstruction from a collection of rich samples.

Definition 10. For a closed subcomplex Y of X, let $\mathscr{F}^Y$ be the sheaf whose stalks are the stalks of $\mathscr{F}$ on Y and zero elsewhere, and whose restrictions are either those of $\mathscr{F}$ on Y or zero as appropriate. There is a surjective sheaf morphism $\mathscr{F} \to \mathscr{F}^Y$ and an induced ambiguity sheaf $\mathscr{F}_Y$ which can be constructed in exactly the same way as $\mathscr{A}$ before.

Thus, the dimension of each stalk of $\mathscr{F}^Y$ is larger than that of any sampling sheaf supported on Y, and the dimension of stalks of $\mathscr{F}_Y$ are therefore as small as or smaller than that of any ambiguity sheaf. Because global sections are determined by their values at the vertices (Proposition 1), obtaining rich samples from $\mathscr{F}^Y$ at all vertices evidently allows reconstruction. This idea works for all degrees of cohomology, which generalizes the notion of oversampling.

Proposition 3. *(Oversampling theorem) If X^k is the closed subcomplex generated by the k-faces of X, then $H^k(X^{k+1};\mathscr{F}_{X^k}) = 0$.*

On the other hand, not taking enough samples leads to an ambiguous reconstruction problem. This can be detected by the presence of nontrivial global sections of the ambiguity sheaf.

Theorem 1. *(Sampling obstruction theorem) Suppose that Y is a closed subcomplex of X and $s : \mathscr{F} \to \mathscr{S}$ is a sampling of sheaves on X supported on Y. If $H^0(X, \mathscr{F}_Y) \neq 0$, then the induced map $H^0(X;\mathscr{F}) \to H^0(X;\mathscr{S})$ is not injective.*

Succinctly, $H^0(X, \mathscr{F}_Y)$ is an obstruction to the recovery of global sections of $\mathscr{F}$ from its samples.

4.1 Proofs of the general sampling theorems

Proof. (of Proposition 3) By direct computation, the k-cochains of $\mathscr{F}_{X^k}$ are

$$\begin{aligned} C^k(X^{k+1};\mathscr{F}_{X^k}) &= C^k(X^{k+1};\mathscr{F})/C^k(X^k;\mathscr{F}) \\ &= \bigoplus_{a \text{ a } k\text{-face of} X} \mathscr{F}(a) / \bigoplus_{a \text{ a } k\text{-face of} X} \mathscr{F}(a) \\ &= 0. \end{aligned}$$

□

As an immediate consequence, $H^0(X;\mathscr{F}_Y) = 0$ when Y is the set of vertices of X.

Proof. (of Theorem 1) We begin by constructing the ambiguity sheaf $\mathscr{A}$ as before so that

$$0 \to \mathscr{A} \to \mathscr{F} \xrightarrow{s} \mathscr{S} \to 0$$

is a short exact sequence of sheaves. Observe that $\mathscr{S} \to \mathscr{F}^Y$ can be chosen to be injective, because the stalks of $\mathscr{S}$ have dimension not more than the dimension of $\mathscr{F}$ (and hence $\mathscr{F}^Y$ also). Thus the induced map $H^0(X;\mathscr{S}) \to H^0(X;\mathscr{F}^Y)$ is also injective. Therefore, by a diagram chase on

$$\begin{array}{ccccc} 0 \to H^0(X;\mathscr{A}) & \longrightarrow & H^0(X;\mathscr{F}) & \xrightarrow{s} & H^0(X;\mathscr{S}) \\ & & \downarrow \cong & & \downarrow \\ 0 \to H^0(X;\mathscr{F}_Y) & \longrightarrow & H^0(X;\mathscr{F}) & \longrightarrow & H^0(X;\mathscr{F}^Y) \end{array}$$

we infer that there is a surjection $H^0(X;\mathscr{A}) \to H^0(X;\mathscr{F}_Y)$. By hypothesis, this means that $H^0(X;\mathscr{A}) \neq 0$, so in particular $H^0(X;\mathscr{F}) \to H^0(X;\mathscr{S})$ cannot be injective. □

5 Examples

This section shows the unifying power of a sheaf-theoretic approach to sampling, by focusing on three rather different examples. The examples differ in terms of how "local" the reconstruction is; those that are less local show a greater impact of the topology of the base space on reconstruction. Specifically, we examine

1. *Bandlimited functions*, in which reconstruction is global. Topology strongly impacts the number of samples required: if we instead consider bandlimited functions on a compact space, we obtain finite Fourier series. (The sampling *rate* is unchanged, however.)
2. *Quantum graphs*, in which reconstruction is somewhat local. Sometimes nontrivial topology in the domain is detected, sometimes not.
3. *Splines*, in which there are only local constraints on the functions. Topology plays almost no role in the reconstruction of splines from their samples.

Of course, the case of PW_B is rather well-known – but we show that it has a sheaf-theoretic interpretation. In rather stark contrast to the case of PW_B is the vector space consisting of the B-splines associated to a particular knot sequence. The functions in this space are determined via a locally finite, piecewise polynomial partition of unity. Since B-splines are determined locally, it makes sense that reconstructing them from local samples is possible. Importantly, sampling theorems obtained for spaces of B-spline are less sensitive to global topological properties.

Spaces of solutions to linear differential equations are a kind of intermediate between PW_B and the space of B-splines. While a degree k differential equation defines its solution locally, there are k linearly independent such solutions. Additionally, the topology of the underlying space on which the differential equation is written impacts the process of reconstruction from samples. [23, 24, 27]

The unifying power of sheaf theory means that all of the examples in this section can be treated in the same way, according to the following procedure:

1. *Encode* the function space to be sampled as a sheaf,
2. *Specify* the sampling sheaf and sampling morphism,
3. *Construct* the ambiguity sheaf associated to the sampling morphism,
4. *Compute* the cohomology of the ambiguity sheaf, and
5. *Collect* conditions for perfect reconstruction based on this computation.

5.1 Bandlimited functions

In this section, we prove the traditional form of the Nyquist theorem by showing that appropriate bandlimiting is a sufficient condition for $H^0(X;\mathscr{A}) = 0$, where $\mathscr{A}$ is an ambiguity sheaf, and X is an abstract simplicial complex for the real line $\mathbb{R}$.

We begin by specifying the following 1-dimensional simplicial complex X. Let $X^0 = \mathbb{Z}$ and $X^1 = \{(n,n+1)\}$. We construct the sheaf $\mathscr{C}$ of signals according to

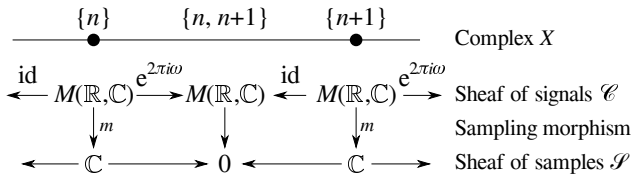

Fig. 7 The sheaves used in proving the traditional Nyquist theorem

their Fourier transforms (see Figure 7) so that for every simplex, the stalk of $\mathscr{C}$ is the vector space $M(\mathbb{R},\mathbb{C})$ of complex-valued measures on $\mathbb{R}$. To relate translation in the spatial domain and the frequency domain, each restriction to the left is chosen to be the identity, and each restriction to the right is chosen to be multiplication by $e^{2\pi i\omega}$. In essence, $\mathscr{C}$ is the sheaf of *local Fourier transforms* of functions on $\mathbb{R}$. Observe that the space of global sections of $\mathscr{C}$ is therefore just $M(\mathbb{R},\mathbb{C})$.

Construct the sampling sheaf $\mathscr{S}$ whose stalk on each vertex is $\mathbb{C}$ and each edge stalk is zero. We construct a sampling morphism by the zero map on each edge, and by the integral

$$m(f) = \int_{-\infty}^{\infty} f(\omega)d\omega = f((-\infty,\infty))$$

on each vertex.

Then the ambiguity sheaf $\mathscr{A}$ has stalks $M(\mathbb{R},\mathbb{C})$ on each edge, and $\{f \in M(\mathbb{R},\mathbb{C}) : m(f) = 0\}$ on each vertex $\{n\}$.

Theorem 2. *(Traditional Nyquist theorem) Suppose we replace $M(\mathbb{R},\mathbb{C})$ with $M_B(\mathbb{R},\mathbb{C})$, the set of measures whose support is contained in $[-B,B]$. Then if $B \leq 1/2$, the resulting ambiguity sheaf $\mathscr{A}_B$ has $H^0(X;\mathscr{A}_B) = 0$. Therefore, each such function can be recovered uniquely from its samples on $\mathbb{Z}$.*

Proof. The elements of $H^0(X;\mathscr{A}_B)$ are given by the measures f supported on $[-B,B]$ for which

$$\int_{-B}^{B} f(\omega)e^{2\pi in\omega}d\omega = (e^{2\pi in\omega}f)([-B,B]) = 0$$

for all n. Observe that if $B \leq 1/2$, this is precisely the statement that the Fourier series coefficients of f all vanish; hence f must vanish. This means that the only global section of $\mathscr{A}_B$ is the zero function. (Ambiguities can arise if $B > 1/2$, because the set of functions $\{e^{-2\pi in\omega}\}_{n\in\mathbb{Z}}$ is then *not* complete.) □

Sampling on the circle can be addressed by a related construction of a sheaf $\mathscr{C}$. As indicated in Figure 8, the stalk over each edge and vertex is still $M(\mathbb{R},\mathbb{C})$. Again, the restrictions are chosen so that left-going restrictions are identities, and the right-going restrictions consist of multiplying measures by $e^{2\pi i\omega}$. Intuitively, this means that functions that are local to an edge or a vertex do not reflect any nontrivial topology.

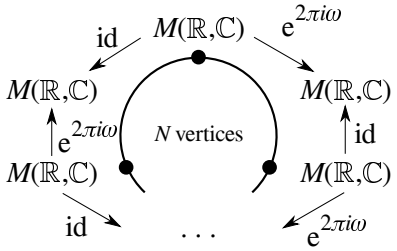

Fig. 8 The sheaf $\mathscr{C}$ of local Fourier transforms of functions on a circle with N vertices

Since the topology is no longer that of a line, there are some important consequences. The space of global sections of $\mathscr{C}$ on the circle is *not* $M(\mathbb{R},\mathbb{C})$, and now depends on the number N of vertices on the circle.[1] One may conclude from a direct computation that the value of any global section at a vertex must be a measure f satisfying

$$e^{2\pi i\omega N} f(\omega) = f(\omega).$$

Of course, this means that the support of f must be no larger than the set of fractions $\frac{1}{N}\mathbb{Z}$ because at each ω either $f(\omega)$ must vanish, or $(e^{2\pi i\omega N} - 1)$ must vanish. Hence a global section describes a function whose (global) Fourier transform is discrete. If we instead consider the sheaf of bandlimited functions $\mathscr{C}_B$ by replacing each $M(\mathbb{R},\mathbb{C})$ by $M_B(\mathbb{R},\mathbb{C})$, the resulting space of global sections is finite dimensional. Perhaps surprisingly, this does not impact the required sampling rate.

Corollary 2. *If $\mathscr{C}_B$ is sampled at each each vertex, then a sampling morphism will fail to be injective on global sections if $B > 1/2$.*

(If $B < 1/2$, some sampling morphisms – such as the zero morphism – are still not injective.)

Proof. Merely observe that for a sampling sheaf $\mathscr{S}$ a necessary condition for injectivity is that

$$\begin{aligned} \dim H^0(\mathscr{C}_B) &\leq \dim H^0(\mathscr{S}) \\ (2N+1)B &\leq N \\ B &\leq N/(2N+1) < 1/2 \end{aligned}$$

□

[1] N must be at least 3 to use an abstract simplicial complex model of the circle. If N is 1 or 2, one must instead use a CW complex. This does not change the analysis presented here.

5.2 Wave propagation (quantum graphs)

A rich source of interesting sheaves arise in the context of differential equations [11]. Sampling problems are interesting in spaces of solutions to differential equations, because they are restricted enough to have relatively relaxed sampling rates. Although a differential equation describes a function locally, continuity and boundary conditions allow topology to influence which of these locally defined functions can be extended globally.

Consider the differential equation

$$\frac{\partial^2 u}{\partial x^2} + k^2 u = 0 \tag{2}$$

on the real line, in which k is a complex scalar parameter called the *wavenumber*. The general solution to this differential equation is the linear combination of two traveling waves, namely

$$u(x) = c_1 e^{ikx} + c_2 e^{-ikx}.$$

This means that locally and globally, a given solution is described by an element of $\mathbb{C}^2$.

Let us now generalize to the case of solutions to (2) over a graph X. In order for the space of solutions to be well defined, it is necessary to assign a length to each edge. We shall write $L(e)$ for the length of edge e, which is a positive real number.

Definition 11. The *transmission line sheaf* $\mathscr{T}$ on X has stalks given by

1. $\mathscr{T}(e) = \mathbb{C}^2$ over each edge e, and
2. $\mathscr{T}(v) = \mathbb{C}^{\deg v}$ over each vertex v, whose degree is $\deg v$.

Without loss of generality, assume that each edge e that is attached to a vertex v is assigned a number from $\{1,\ldots,\deg v\}$. If e is the m-th edge attached to vertex v, the restriction $\mathscr{T}(v \rightsquigarrow e_m)$ is given by

$$\mathscr{T}(v \rightsquigarrow e_m)(u_1,\ldots,u_{\deg v}) = \begin{cases} \left(u_m, e^{-ikL(e_m)}\left(\frac{2}{\deg v}\sum_{j=1}^{\deg v} u_j - u_m\right)\right) & \text{if } [e_m : v] = 1 \\ \left(e^{ikL(e_m)}\left(\frac{2}{\deg v}\sum_{j=1}^{\deg v} u_j - u_m\right), u_m\right) & \text{if } [e_m : v] = -1 \end{cases}$$

Although the definition of a transmission line sheaf is combinatorial, it is an accurate model of the space of solutions on its *geometric realization*. Under an appropriate definition of the Laplacian operator on the geometric realization of X (such as is given in [40, 2, 1]), the sheaf of solutions to (2) is isomorphic to a transmission line sheaf. [27] Additionally, there are sensible definitions for bandlimitedness in this geometric realization, which give rise to Shannon-Nyquist theorems. [23, 24] However, our focus will remain combinatorial and topological. We note that others [25, 26] have obtained results in general combinatorial settings, though we will focus on the impact of the topology of X on sampling requirements.

The easiest sampling result for transmission line sheaves follows immediately from Proposition 1, namely that a global section of a transmission line sheaf $\mathscr{T}$ on X is completely specified by its values on the vertices of X.

This result is clearly inefficient; merely consider the simplicial complex X for the real line with vertices $X^0 = \mathbb{Z}$ and edges $X^1 = \{(n, n+1)\}$. We have already seen that the space of sections of a tranmission line sheaf on X is merely $\mathbb{C}^2$, yet oversampling would have us collect samples at infinitely many vertices! The missing insight is that the topology of X impacts the global sections of a transmission line sheaf.

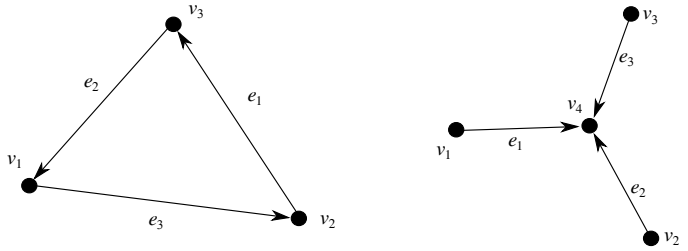

Fig. 9 The graph X for Example 11 (left) and Y for Example 10 (right)

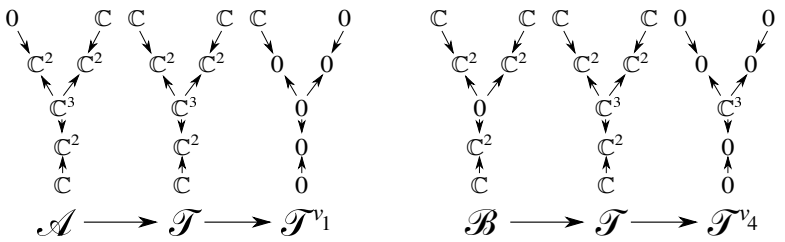

Fig. 10 Sampling a star at a leaf (left) and at its center (right)

Changing the edge length in the simplicial complex model for $\mathbb{R}$ does not change the space of global sections of a transmission line sheaf. Another siuation in which edge length does not matter is shown in the next example.

Example 10. Consider the space Y shown at right in Figure 9. The coboundary map d^0 for a transmission line sheaf over Y is given by

$$\begin{pmatrix}
-e^{ikL(e_1)} & 0 & 0 & 1 & 0 & 0 \\
-1 & 0 & 0 & -\frac{1}{3}e^{-ikL(e_1)} & \frac{2}{3}e^{-ikL(e_1)} & \frac{2}{3}e^{-ikL(e_1)} \\
0 & -e^{ikL(e_2)} & 0 & 0 & 1 & 0 \\
0 & -1 & 0 & \frac{2}{3}e^{-ikL(e_2)} & -\frac{1}{3}e^{-ikL(e_2)} & \frac{2}{3}e^{-ikL(e_2)} \\
0 & 0 & -e^{ikL(e_3)} & 0 & 0 & 1 \\
0 & 0 & -1 & \frac{2}{3}e^{-ikL(e_3)} & \frac{2}{3}e^{-ikL(e_3)} & -\frac{1}{3}e^{-ikL(e_3)}
\end{pmatrix},$$

which has rank 5 – in other words $\dim H^0(Y;\mathscr{T}) = 1$ for any transmission sheaf $\mathscr{T}$. This means that reconstruction of sections requires at least one dimension of measurements; a lower bound. If we consider a sampling morphism $\mathscr{T} \to \mathscr{T}^Z$ for some set of vertices Z, this induces an injective map on global sections. To see this, consider sampling at any one of the leaf nodes or at the center.

If we sample at a leaf node only, as shown in Figure 10 at left, the ambiguity sheaf $\mathscr{A}$ has a coboundary matrix given by

$$\begin{pmatrix}
0 & 0 & 1 & 0 & 0 \\
0 & 0 & -\frac{1}{3}e^{-ikL(e_1)} & \frac{2}{3}e^{-ikL(e_1)} & \frac{2}{3}e^{-ikL(e_1)} \\
-e^{ikL(e_2)} & 0 & 0 & 1 & 0 \\
-1 & 0 & \frac{2}{3}e^{-ikL(e_2)} & -\frac{1}{3}e^{-ikL(e_2)} & \frac{2}{3}e^{-ikL(e_2)} \\
0 & -e^{ikL(e_3)} & 0 & 0 & 1 \\
0 & -1 & \frac{2}{3}e^{-ikL(e_3)} & \frac{2}{3}e^{-ikL(e_3)} & -\frac{1}{3}e^{-ikL(e_3)}
\end{pmatrix},$$

which has rank 5, implying that the ambiguity sheaf has only trivial global sections.

On the other hand, sampling at the center yields a different ambiguity sheaf $\mathscr{B}$, whose coboundary matrix is

$$\begin{pmatrix}
-e^{ikL(e_1)} & 0 & 0 \\
-1 & 0 & 0 \\
0 & -e^{ikL(e_2)} & 0 \\
0 & -1 & 0 \\
0 & 0 & -e^{ikL(e_3)} \\
0 & 0 & -1
\end{pmatrix},$$

which also has trivial kernel.

If we instead consider a different topology, for instance a circle, then edge lengths do have an impact on the global sections of the resulting transmission line sheaf.

Example 11. Consider the simplicial complex X shown at left in Figure 9, in which the edges are oriented as marked. Because X has a nontrivial loop, the lengths of the edges impact the space of global sections of a transmission line sheaf over X. Specifically, if $\mathscr{T}$ is a transmission line sheaf, its coboundary $d^0 : C^0(X;\mathscr{T}) \to C^1(X;\mathscr{T})$ is given by

$$\begin{pmatrix} 0 & 0 & -e^{ikL(e_1)} & 0 & 1 & 0 \\ 0 & 0 & 0 & -1 & 0 & e^{-ikL(e_1)} \\ 1 & 0 & 0 & 0 & -e^{ikL(e_2)} & 0 \\ 0 & e^{-ikL(e_2)} & 0 & 0 & 0 & -1 \\ -e^{ikL(e_3)} & 0 & 1 & 0 & 0 & 0 \\ 0 & -1 & 0 & e^{-ikL(e_3)} & 0 & 0 \end{pmatrix}.$$

This matrix has full rank unless $e^{-ik(L(e_1)+L(e_2)+L(e_3))} = 1$, a condition called *resonance*. Therefore, the space of global sections of $\mathscr{T}$ has dimension

$$\dim H^0(X;\mathscr{T}) = \begin{cases} 2 & \text{if } k(L(e_1)+L(e_2)+L(e_3)) \in 2\pi\mathbb{Z} \\ 0 & \text{otherwise} \end{cases},$$

and an easy calculation shows that sampling at any one of the vertices results in an injective map on global sections.

Based on the previous examples, a sound procedure is to consider the dimension of the space of global sections of $\mathscr{T}$ to be a lower bound on how much information is to be obtained through sampling. (Clearly, this may not be enough in some situations, especially if the sampling morphisms are not injective on stalks.) As described in [27], a general lower bound on the dimension of $H^0(\mathscr{T})$ is $n+1$, where n is the number of resonant loops. (A tighter lower bound exists, but its expression is complicated by the presence of degree 1 vertices.) Therefore, topology plays an important role in acquiring enough information to recover global sections of $\mathscr{T}$ from samples.

5.3 Polynomial splines

Section 5.1 showed how limiting the support of the Fourier transform of a function permitted it to be reconstructed by its values at a discrete subset. Because of the Paley-Weiner theorem, the smoothness of a function is reflected in the decay of its Fourier transform. On the other hand, Section 5.2 showed that applying smoothness constraints directly to the function also enables perfect recovery. This suggests that as we consider smoother functions, we can reconstruct them from more widely spaced samples.

In this section, we consider sampling from polynomial splines, which are functions whose smoothness is explicitly controlled. A C^k degree n polynomial spline has k continuous derivatives and is constructed piecewise from degree n polynomial segments (see Figure 11). Because of this, a polynomial spline is infinitely differentiable on all of its domain except at a discrete set of *knot points*, where it has k continuous derivatives.

For instance, consider a degree n polynomial spline that has two knots: one at 0 and one at L. Require it to have C^{n-1} smoothness across its three segments: $(-\infty, 0)$,

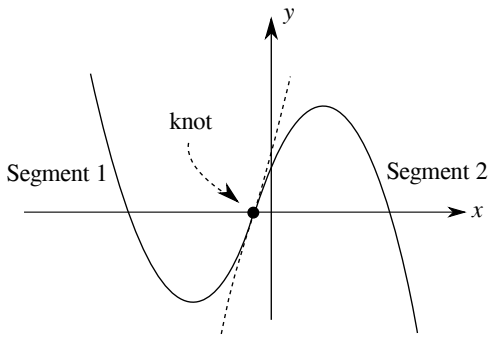

Fig. 11 A polynomial spline with two quadratic segments, joined at a knot with continuous first derivatives

$(0,L)$, and (L,∞). To obtain $n-1$ continuous derivatives at $x=0$, such a spline should have the form

$$f(x) = \begin{cases} a_n^- x^n + \sum_{k=0}^{n-1} a_k x^k & \text{for } x \le 0 \\ a_n^+ x^n + \sum_{k=0}^{n-1} a_k x^k & \text{for } 0 \le x \le L \end{cases}$$

In a similar way, to obtain $n-1$ continuous derivatives at $x=L$, the spline should be of the form

$$f(x) = \begin{cases} b_n^- (x-L)^n + \sum_{k=0}^{n-1} b_k (x-L)^k & \text{for } 0 \le x \le L \\ b_n^+ (x-L)^n + \sum_{k=0}^{n-1} b_k (x-L)^k & \text{for } x \ge L \end{cases}$$

But clearly on $0 \le x \le L$, these two definitions should agree so that f is a well-defined function. This means that for all x,

$$\begin{aligned} a_n^+ x^n + \sum_{k=0}^{n-1} a_k x^k &= b_n^- (x-L)^n + \sum_{k=0}^{n-1} b_k (x-L)^k \\ &= \sum_{i=0}^{n-1} b_i \sum_{k=0}^{i} \binom{i}{k} x^k (-L)^{i-k} + b_n^- \sum_{k=0}^{n} \binom{n}{k} x^k (-L)^{n-k} \\ &= \sum_{k=0}^{n-1} x^k \sum_{i=k}^{n-1} \binom{i}{k} (-L)^{i-k} b_i + \sum_{k=0}^{n} x^k \binom{n}{k} (-L)^{n-k} b_n^- \\ &= \sum_{k=0}^{n-1} x^k \left(\sum_{i=k}^{n-1} \binom{i}{k} (-L)^{i-k} b_i + \binom{n}{k} (-L)^{n-k} b_n^- \right) + b_n^- x^n. \end{aligned}$$

By linear independence, this means that

$$a_n^+ = b_n^-$$
$$a_k = \sum_{i=k}^{n-1} \binom{i}{k}(-L)^{i-k}b_i + \binom{n}{k}(-L)^{n-k}b_n^-$$

Notice that if the a variables are given, then this is a triangular system for the b variables.

Using this computation, we can define $\mathscr{PS}^n$, the sheaf of n-degree polynomial splines on $\mathbb{R}$ with knots at each of the integers. This sheaf is built on the simplicial complex X for $\mathbb{R}$, whose vertices are $X^0 = \mathbb{Z}$ and whose edges are $X^1 = \{(m, m+1)\}$. Because a degree n spline at any given knot can be defined by $n+2$ real values on the two segments adjacent to that knot, we assign

$$\mathscr{PS}^n(\{m\}) = \mathbb{R}^{n+2}$$

for each $m \in \mathbb{Z}$. Specifically, we will think of these as defining $(a_0, \ldots, a_{n-1}, a_n^-, a_n^+)$ in our calculation above. The spline on each segment is merely a degree n polynomial, so that

$$\mathscr{PS}^n(\{(m, m+1)\}) = \mathbb{R}^{n+1}.$$

For each knot, there are two restriction maps: one to the left and one to the right. They are given by

$$\mathscr{PS}^n(\{m\} \rightsquigarrow \{m, m+1\}) = \begin{pmatrix} 1 & \cdots & 0 & 0 & 0 \\ \vdots & & \vdots & \vdots & \vdots \\ 0 & \cdots & 1 & 0 & 0 \\ 0 & \cdots & 0 & 0 & 1 \end{pmatrix}$$

and

$$\mathscr{PS}^n(\{m+1\} \rightsquigarrow \{m, m+1\}) = \begin{pmatrix} 1 & -L & L^2 & -L^3 & \cdots & (-L)^n & 0 \\ & 1 & -2L & 3L^2 & \cdots & \binom{n}{1}(-L)^{n-1} & 0 \\ & & 1 & -3L & \cdots & \binom{n}{2}(-L)^{n-2} & 0 \\ & & & 1 & \cdots & \binom{n}{3}(-L)^{n-3} & 0 \\ & & & & & \vdots & \vdots \\ & & & & & \binom{n}{k}(-L)^{n-k} & 0 \\ & & & & & \vdots & \vdots \\ & & & & & 1 & 0 \end{pmatrix}$$

Remark 5. Observe that $L = 1$, $\mathscr{PS}^1$ reduces to the sheaf $\mathscr{PL}$ given in Example 7, for the special case of the graph being a line (so all vertices have degree 2).

Lemma 1. *Consider $\mathscr{PS}^n$ on the simplicial complex shown in Figure 12, which has $k+2$ vertices and and $k+1$ edges. The sheaf has nontrivial global sections that*

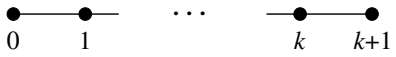

Fig. 12 The simplicial complex used in Lemma 1

vanish at the endpoints if and only if $k > n+1$. If $k \leq n+1$, then the only global section which vanishes at both endpoints is the zero section.

Proof. Consider the ambiguity sheaf $\mathscr{A}$ associated to the sampling morphism $\mathscr{P}\mathscr{S}^n \to (\mathscr{P}\mathscr{S}^n)^Y$ where Y consists of the two endpoints. Observe that the global sections of $\mathscr{A}$ correspond to global sections of $\mathscr{P}\mathscr{S}^n$ that vanish at the endpints, so the lemma follows by reasoning about $H^0(\mathscr{A})$. The matrix for the coboundary map $d^0 : C^0(\mathscr{A}) \to C^1(\mathscr{A})$ has a block structure

$$\begin{pmatrix} A & 0 & & 0 \\ B & A & \cdots & 0 \\ 0 & B & & 0 \\ & & & \vdots \\ & & & A \\ & & & B \end{pmatrix},$$

where the $(n+1) \times (n+2)$ blocks are

$$A = \begin{pmatrix} 1 & \cdots & 0 & 0 & 0 \\ \vdots & & \vdots & \vdots & \vdots \\ 0 & \cdots & 1 & 0 & 0 \\ 0 & \cdots & 0 & 0 & 1 \end{pmatrix}, \; B = \begin{pmatrix} -1 & \cdots & * & * & 0 \\ \vdots & & \vdots & \vdots & \vdots \\ 0 & \cdots & -1 & * & 0 \\ 0 & \cdots & 0 & -1 & 0 \end{pmatrix}.$$

Clearly both such blocks are of full rank. Thus the coboundary matrix has a nontrivial kernel whenever it has more rows than columns:

$$\begin{aligned} (n+2)k &> (n+1)(k+1) \\ nk+1k &> nk+n+k+1 \\ k &> n+1 \end{aligned}$$

as desired. □

This lemma implies that unambiguous reconstruction from samples is possible provided the gaps between samples are small enough. Increased smoothness allows the gaps to larger without inhibiting reconstruction. Because of this, it is convenient to define distances between vertices.

Definition 12. On a graph G, define the *edge distance* between two vertices v, w to be

$$\mathrm{ed}(v,w) = \begin{cases} \min_p\{\#\text{ edges in } p \text{ such that } p \text{ is a} \\ \text{PL-continuous path from } v \to w\} \\ \infty \text{ if no such path exists} \end{cases}$$

From this, the maximal distance to a vertex set Y is

$$\mathrm{med}(Y) = \max_{x \in X^0}\{\min_{y \in Y} \mathrm{ed}(x,y)\}.$$

Corollary 3. *Suppose that $Y \subseteq \mathbb{Z}$, which we take to be a subset of vertices of X. If $\mathrm{med} Y \leq n+2$, then the sampling morphism $\mathscr{PS}^n \to (\mathscr{PS}^n)^Y$ induces an injective linear map on global sections.*

Observe that if we instead consider sampling of polynomial splines on the circle, very little changes. In particular, the proof of Lemma 1 doesn't change at all. Indeed, we can change the context from polynomial splines on the real line to piecewise linear functions $\mathscr{PL}$ on a graph. We'll focus on the special case of a sampling morphism $s : \mathscr{PL} \to \mathscr{PL}^Y$ where Y is a subset of the vertices of X. Excluding one or two vertices from Y does not prevent reconstruction in this case, because the samples include information about slopes along adjacent edges.

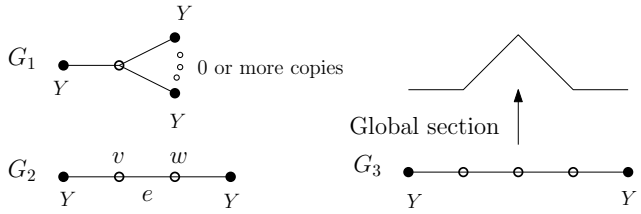

Fig. 13 Graphs G_1, and G_2 (left) and G_3 (right) for Lemma 2. Filled vertices represent elements of Y, empty ones are in the complement of Y.

Lemma 2. *Consider $\mathscr{PL}_Y$, the subsheaf of $\mathscr{PL}$ whose sections vanish on a vertex set Y and the graphs G_1, G_2, and G_3 as shown in Figure 13. There are no nontrivial sections of $\mathscr{PL}_Y$ on G_1 and G_2, but there are nontrivial sections of $\mathscr{PL}_Y$ on G_3.*

Proof. If a section of $\mathscr{PL}$ vanishes at a vertex x with degree n, this means that the value of the section there is an $(n+1)$-dimensional zero vector. The value of the section on every edge adjacent to x is then the 2-dimensional zero vector. Since the dimensions in each stalk of $\mathscr{PL}$ represent the value of the piecewise linear function and its slopes, linear extrapolation to the center vertex in G_1 implies that its value is zero too.

A similar idea applies in the case of G_2. The stalk at v has dimension 3. Any section at v that extends to the left must actually lie in the subspace spanned by $(0,0,1)$ (coordinates represent the value, left slope, right slope respectively). In the same way, any section at w that extends to the right must lie in the subspace spanned

by $(0,1,0)$. Any global section must extend to e, which must therefore have zero slope and zero value.

Finally G_3 has nontrivial global sections, spanned by the one shown in Figure 13. □

Proposition 4. *(Unambiguous sampling) Consider the sheaf $\mathscr{PL}$ on a graph X and $Y \subseteq X^0$. Then $H^0(X;\mathscr{F}_Y) = 0$ if and only if med$(Y) \leq 1$.*

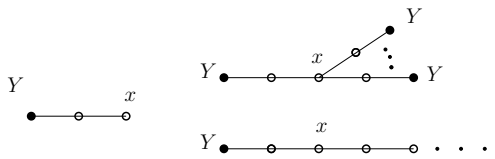

Fig. 14 The three families of subgraphs that arise when med$(Y) > 1$. Filled vertices represent elements of Y, empty ones are in the complement of Y.

Proof. ($\Leftarrow$) Suppose that $x \in X^0 \backslash Y$ is a vertex not in Y. Then there exists a path with one edge connecting it to Y. Whence we are in the case of G_1 of Lemma 2, so any section at x must vanish.

($\Rightarrow$) By contradiction. Assume med$(Y) > 1$. Without loss of generality, consider $x \in X^0 \backslash Y$, whose distance to Y is exactly 2. Then one of the subgraphs shown in Figure 14 must be present in X. But case G_3 of Lemma 2 makes it each of these has nontrivial sections at x, merely looking at sections over the subgraph. □

Proposition 5. *(Non-redundant sampling) Consider the case of $s : \mathscr{PL} \to \mathscr{PL}^Y$. If $Y = X^0$, then $H^1(X;\mathscr{A}) \neq 0$. If Y is such that med$(Y) \leq 1$ and $|X^0 \backslash Y| + \sum_{y \notin Y} \deg y = 2|X^1|$, then $H^1(X;\mathscr{A}) = 0$.*

Proof. The stalk of $\mathscr{A}$ over each edge is $\mathbb{R}^2$, and the stalk over a vertex in Y is trival. However, the stalk over a vertex of degree n not in Y is $\mathbb{R}^{n+1}$. Observe that if $H^0(X;\mathscr{A}) = 0$, then $H^1(X;\mathscr{A}) = C^1(X;\mathscr{A})/C^0(X;\mathscr{A})$. Using the degree sum formula in graph theory, we compute that $H^1(X;\mathscr{A})$ has dimension $2|X^1| - \sum_{y \notin Y}(\deg y + 1)$. □

6 Conclusions

This chapter has shown that exact sequences of sheaves are a unifying principle for sampling theory. These tools provide a general basis for discussing locality, and reveal general, precise conditions under which reconstruction succeeds. Several sampling problems for bandlimited and non-bandlimited functions were discussed. The use of sheaves in sampling is essentially unexplored otherwise, and there remain

many open questions. We discuss two such questions here with a relationship to bandlimited functions.

The first question is rather connected to the existing literature on sampling. We showed that although the number of samples required for the reconstruction of a bandlimited function may vary, the sampling *rate* necessary appears to be the same. Specifically, we found that the bandwidth B required for reconstructing functions on the real line and the circle was constrained to be less than $1/2$. Others (for instance [23, 26]) have found that this remains the case for other domains as well. Therefore, it seems fitting to ask "Is the $B < 1/2$ in the Nyquist Theorem invariant with respect to changes in topology and geometry?"

The second question is a bit more subtle. If a function is bandlimited, this means that its Fourier transform has bounded support. However, the Fourier transform is intimately related to the spectrum of the Laplacian operator. On the other hand, because cohomological obstructions play a role in sampling, we had to consider the cochain complex for sheaves of functions. These two threads of study are in fact closely related through Hodge theory via the study of Hilbert complexes [5]. Indeed, if $\mathscr{F}$ and $\mathscr{S}$ are sheaves of Hilbert spaces then their cochain complexes are Hilbert complexes, and they have a Hodge decomposition, along with an associated Laplacian operator. Therefore, there may be a way to discuss the concept of "bandwidth" for general sheaves of Hilbert spaces.

Acknowledgements This work was partly supported under Federal Contract No. FA9550-09-1-0643. The author also wishes to thank the editor for the invitation to write this chapter. Portions of this article appeared in the proceedings of SampTA 2013, published by EURASIP.

References

1. Mathhew Baker and Robert Rumely. Harmonic analysis on metrized graphs. *Canad. J. Math.*, 59(2):225–275, 2007.
2. Matthew Baker and Xander Faber. Metrized graphs, Laplacian operators, and electrical networks. In *Quantum graphs and their applications*, pages 15–34, 2006.
3. J. Benedetto and W. Heller. Irregular sampling and the theory of frames: I. *Note di Mathematica*, 10(1):103–125, 1990.
4. Glen Bredon. *Sheaf theory*. Springer, 1997.
5. J. Brüning and M. Lesch. Hilbert complexes. *Journal of Functional Analysis*, 108:88–132, 1992.
6. F. Chazal, D. Cohen-Steiner, and A. Lieutier. A sampling theory for compact sets in euclidean space. *Discrete Comput. Geom.*, 41:461–479, 2009.
7. J. Curry. Sheaves, cosheaves and applications, `arxiv:1303.3255`. 2013.
8. J. Curry, R. Ghrist, and M. Robinson. Euler calculus and its applications to signals and sensing. In Afra Zomorodian, editor, *Proceedings of Symposia in Applied Mathematics: Advances in Applied and Computational Topology*, 2012.
9. P. Dragotti, M. Vetterli, and T. Blue. Sampling moments and reconstructing signals of finite rate of innovation: Shannon meets Strang-Fix. *IEEE Trans. Sig. Proc.*, 55(5), May 2007.
10. M. Ebata, M. Eguchi, S. Koizumi, and K. Kumahara. Analogues of sampling theorems for some homogeneous spaces. *Hiroshima Math. J.*, 36:125–140, 2006.

11. L. Ehrenpreis. Sheaves and differential equations. *Proc. AMS*, 7(6):1131–1138, December 1956.
12. H. Feichtinger and I. Pesenson. A reconstruction method for band-limited signals on the hyperbolic plane. *Sampl. Theory Signal Image Process.*, 4(2):107–119, 2005.
13. H.G. Feichtinger and K. Gröchenig. Theory and practice of irregular sampling. *Wavelets: mathematics and applications*, pages 305–363, 1994.
14. R. Ghrist and Y. Hiraoka. Applications of sheaf cohomology and exact sequences to network coding. *preprint*, 2011.
15. R. Godement. *Topologie algebrique et théorie des faisceaux*. Herman, Paris, 1958.
16. K. Gröchening. Reconstruction algorithms in irregular sampling. *Mathematics of Computation*, 59(199):181–194, July 1992.
17. A. Hatcher. *Algebraic Topology*. Cambridge University Press, 2002.
18. John H. Hubbard. *Teichmüller Theory, volume 1*. Matrix Editions, 2006.
19. B. Iverson. *Cohomology of Sheaves*. Aarhus universitet, Matematisk institut, 1984.
20. J. Lilius. Sheaf semantics for Petri nets. Technical report, Helsinki University of Technology, Digital Systems Laboratory, 1993.
21. P. Niyogi, S. Smale, and S. Weinberger. Finding the homology of submanifolds with high confidence from random samples. In R. Pollack, J. Pach, and J. E. Goodman, editors, *Twentieth Anniversary Volume*, pages 1–23. Springer New York, 2009.
22. I. Pesenson. Sampling of band-limited vectors. *Journal of Fourier Analysis and Applications*, 7(1):93–100, 2001.
23. I. Pesenson. Band limited functions on quantum graphs. *Proceedings of the American Mathematical Society*, 133(12):3647–3656, 2005.
24. I. Pesenson. Analysis of band-limited functions on quantum graphs. *Applied and Computational Harmonic Analysis*, 21(2):230–244, 2006.
25. I. Pesenson. Sampling in Paley-Wiener spaces on combinatorial graphs. *Transactions of the American Mathematical Society*, 360(10):5603, 2008.
26. I.Z. Pesenson and M.Z. Pesenson. Sampling, filtering and sparse approximations on combinatorial graphs. *Journal of Fourier Analysis and Applications*, 16(6):921–942, 2010.
27. M. Robinson. Inverse problems in geometric graphs using internal measurements, `arxiv:1008.2933`. 2010.
28. M. Robinson. Asynchronous logic circuits and sheaf obstructions. *Electronic Notes in Theoretical Computer Science*, pages 159–177, 2012.
29. M. Robinson. Understanding networks and their behaviors using sheaf theory. In *GlobalSIP*, 2013.
30. M. Robinson. *Topological Signal Processing*. Springer, 2014.
31. A. Schuster and D. Varolin. Interpolation and sampling for generalized Bergman spaces on finite Riemann surfaces. *Revista Matemática Iberoamericana*, 24(2):499–530, 2008.
32. Allen Shepard. *A cellular description of the derived category of a stratified space*. PhD thesis, Brown University, 1980.
33. Allen Shepard. *A cellular description of the derived category of a stratified space*. PhD thesis, Brown University, 1985.
34. S. Smale and D.X. Zhou. Shannon sampling and function reconstruction from point values. *Bulletin of the American Mathematical Society*, 41(3):279–306, 2004.
35. M. Unser. Splines: A perfect fit for signal and image processing. *Signal Processing Magazine, IEEE*, 16(6):22–38, 1999.
36. M. Unser. Sampling–50 years after Shannon. *Proceedings of the IEEE*, 88(4):569–587, 2000.
37. M. Unser and J. Zerubia. A generalized sampling theory without band-limiting constraints. *Circuits and Systems II: Analog and Digital Signal Processing, IEEE Transactions on*, 45(8):959–969, 1998.
38. R.G. Vaughan, N.L. Scott, and D.R. White. The theory of bandpass sampling. *Signal Processing, IEEE Transactions on*, 39(9):1973–1984, 1991.
39. M. Vetterli, P. Marziliano, and T. Blu. Sampling signals with finite rate of innovation. *Signal Processing, IEEE Transactions on*, 50(6):1417–1428, 2002.
40. Shouwu Zhang. Admissible pairing on a curve. *Invent. Math.*, 112(1):171–193, 1993.